Experimental Results of LightSAR Mission Planning Using a Market-Based System

Randii R. Wessen
Jet Propulsion Laboratory
California Institute of Technology

David Porter
Economic Science Laboratory
University of Arizona

Jeffrey Hilland
Jet Propulsion Laboratory
California Institute of Technology

Abstract

The allocation of scarce spacecraft resources to multiple users has always been a difficult process. This difficulty arises from the fact that there are never enough resources (e.g., data volume, integration time, spacecraft power, etc.) to meet the stated requirements of the scientific investigators who compete to acquire their desired data sets. To help solve this problem, a market-based process was developed to allocate on-orbit resources for the LightSAR mission, a joint NASA/commercial endeavor. LightSAR chose to evaluate the utility of a market-based system as part of its mission concept study phase. A market-based system was selected based on its prior successes in allocating resources for the Cassini spacecraft and the Space Shuttle. These experiments demonstrated that a market-based system could provide results comparable to the current methods for allocating resources while requiring a smaller workforce and a shorter period of time to develop.

Introduction

One of the most time-consuming activities performed during mission operations is the conflict-resolution process for determining which Investigator's data requests take precedence over another. These conflicts occur any time there are multiple science objectives for a given instrument or multiple instruments with unique objectives, i.e., when the demands for spacecraft resources outstrip the available supply.

In November 1997, the LightSAR (Lightweight Synthetic Aperture Radar) Pre-Project undertook an investigation into a market-based system to determine if it could be a cost-effective planning tool. Data acquisition planning for the LightSAR radar will be complex because of the anticipated large demand for use of the payload and the complicated joint NASA/Commercial structure of the project.

Typically, Project officials are placed in the difficult position of not being able to give Science Investigators all the resources they request. By resources we mean payload data acquisition time. Knowing this a priori give Investigators the incentive to request more resources than they need since they will not receive all the resources they request.

The usual process for scheduling data acquisition requests on Earth-orbiting radar spacecraft has involved a "committee-driven" approach. This approach requires individual Investigators to submit requests for specific spacecraft resources to a "neutral party," namely the Sequence Integrators. These individuals integrate the requests into a single time-ordered listing of events that do not violate resource constraints. The Integrator's goal is to produce a "conflict-free" listing that maximizes the overall return for the mission while being "fair" to each Investigator. Fair in this context means that every attempt is made to integrate each Investigator's highest ranked requests into the listing.

The approach used by Integrators is sometimes referred to as a "Serial Draft" or "Serial Dictator" method. The Integrator starts with one Investigator and selects his highest ranked request. The Integrator then moves to the next Investigator, selecting that person's highest ranked request. This process continues until the highest ranked request from each Investigator is incorporated into the time-ordered listing. Once completed, the Integrator selects the second highest ranked request. This time, however, the order is reversed. The Investigator who had his highest ranked request selected last, gets the next highest ranked request selected first, and so on. If there are not enough resources left for the current request, then the Integrator selects the next highest ranked request from that Investigator's prioritized list. This selection continues until either all of the requests are implemented or the remaining spacecraft resources can not accommodate any other requests.

Once the Integrator develops a listing, it is presented to the Science Investigators for evaluation. During this evaluation process, comments are submitted. Typically, those Investigators who have their requests incorporated evaluate the listing quite high. Those Investigators whose requests are not realized however, evaluate the listing low. Since there is no direct mechanism to control the amount of appeals, most Investigators appeal to have more data acquisition time.

Appeals involve presenting the merits of one Investigator's request over another to the Project Scientist or some other governing board. This appeal produces what is commonly refereed to as a "Dead Weight Loss."[1] That is, the Investigator who is not awarded additional acquisition time has invested their time and effort for the appeal and have received no return. The end result is that multiple meetings with multiple appeals and re-integrations occur until the time for developing the time-ordered listing has expired.

To reduce the amount of time and workforce needed to produce a conflict-free, time-ordered listing, LightSAR decided to evaluate a market-based system. These systems have been used for centuries in economics and have recently been successfully applied to space missions. For example, a market-based approach was used during Cassini development from 1993 to 1995, to control the science instrument's demand for resources. Results from the Cassini's Resource Exchange (CRE) showed that instrument cost growth was less than 1%, and instrument mass growth decreased by 7%.[2] Prior missions usually had mass and cost growths that exceeded well over 100%.

The CRE system was then transferred to Southern California's Air Quality Management District (AQMD) and is currently being used to successfully control smog emissions in Southern California. Market-based systems have also been successfully tested for the Federal Communications Commission (FCC) Spectrum Auction for Personnel Communication Service Licenses. A prototype system was also developed for manifesting Space Shuttle Secondary Payloads.[3]

For LightSAR, a committee-driven approach was compared to a market-based approach. A market-based system uses "rights" and "trades" to resolve conflicts, instead of educated guesses made by a subjective third party. Each Science Investigator is allocated a "currency" for expressing the relative importance of one request over another. This currency, which we dub "Priority Points," is budgeted to each Investigator who in turn assigns them to their data takes to define the "worth" of the request. Investigators are free to express the relative importance of their requests and make trades among themselves to enhance their positions. A further advantage is that this market-based system resides on the Internet and allows Investigators around the world to develop a resource timeline remotely.

Testbed Experiments

In order to test the ability of such a system to develop an efficient timeline of data takes, a set of controlled laboratory experiments was conducted. The use of experiments to evaluate comparative allocation systems has been a reliable source of scientific data. The methodology of experimental economics is similar to the use of wind tunnels to test airfoil designs.

The main components of an experiment are: 1) defining what is to be allocated; 2) setting individual incentives; and, 3) defining the process by which resources are allocated. For this experiment, we defined fixed-duration data take requests as the resources to be allocated. There were four data takes per orbit and four orbits per planning period. We used two planning periods where individuals could carry forward any unused Priority Points from Period One to Period Two.

To set up the experiment, Investigators from the radar community were asked to define the value of each request as a function of its scientific or commercial objectives. Table 1 shows an example of how an Investigator might define the value of the data take requests.

Location	Orbit Number	Data Take Number	Rank	Mission Value
Vietnam	1	1	1	60
Kuala Lumpur	1	3	2	45
Indonesia	1	4	2	35
Cambodia	2	3	3	10

Table 1. An example of Dual Polarimetry Data Take requests.

In this example, the Dual Polarimetry Investigator ranked data takes and then assigned them a mission value. Notice that in this example, the

Investigator ranked data takes for Kuala Lumpur and Indonesia with a rank of two. If this information alone were given to the Integrator as usual, they would assume each location was equally important and would assign the data take easiest to incorporate into the time-ordered listing. However, Kuala Lumpur had a value of 45 while Indonesia had only 35. The two locations were not equal, so the Dual Polarimetry Investigator in reality did have a preference.

Defining a mission value for each data take has another advantage over a simple ranking. It provides tradeoff information and the relative worth of each data take. As an example, using Table 1, an Integrator would try to incorporate Vietnam (rank=1), followed by either Kuala Lumpur (rank=2) or Indonesia (rank=2). However, using the value column, the Dual Polarimetry Investigation would produce a greater mission value if Kuala Lumpur and Indonesia (45+35=80 points) could be incorporated into the timeline instead of the number one ranked Vietnam (60 points) request. This example shows that a simple ranking does not provide enough information to produce the highest value time-ordered listing.

For the LightSAR experiment, undergraduates from the California Institute of Technology were used as test Investigators. The students' incentive was financial compensation as they were paid as a function of how well they were able to get their data takes into the time-ordered listing. Students were assigned one of the five roles: Dual Polarimetry, Quad Polarimetry, Interferometry, ScanSAR, Hi-Res Strip, or Spotlight. Students then bid for particular data takes that would provide the highest values. A typical student's bid might look as indicated in Table 2.

Status	Location	Orbit Number	Data Take Number	Bid (Priority Points)
New	Vietnam	1	1	25

Table 2. A typical bid showing its status and the number of Priority Points.

A bid is simply an expression of the level of importance a particular location has to the student Investigator. The higher the number of Priority Points bid, the greater the value of the request to the student. For the experiment, the bidding proceeded in rounds. Once submitted, successful bids could not be retracted. This rule ensured that bids were monotonic and that the process would convergence.

Once bids were received from each student, the round was closed and a solution that maximized the sum of Priority Points bid and which produced a conflict-free schedule was computed. Once solved, the next round began. Students could then see if their data take requests were incorporated into the listing or determine the number of Priority Points needed to "out bid" another user's successful data take request. The students had the choice to resubmit their bid with a larger number of points or choose some other data take. Once again, when all bids were received, the round was closed and then solved for the greatest point value. The rounds continued until the value of the time-ordered listing did not increase by 10% of the value of the previous round.

Rounds lasted approximately 5 minutes apiece, allowing the experimenters to run many experiments in a relatively short period of time. The sheer number of resulting iterations allowed the students to validate the experiment's design, find flaws in the operations, and to vary initial conditions.

Once student experiments were complete, experiments were performed with the science community. In these cases, rounds were much longer with one round in the morning and one in the afternoon. A Science Investigator could log-on to the LightSAR experiment website, evaluate the time-ordered listing, submit bids, and then log-off. The conditions for ending the planning period were the same as for the student experiments. That is, the planning period ended when the value of the time-ordered listing did not increase by 10% of the value of the previous round.

One interesting problem was determining how to end an experiment (i.e., the planning period). If not chosen careful, a poorly designed ending could produce undesirable results. For example, a specific time can be used for the close of a planning period. However, this method produces the undesirable incentive for all Investigators to wait until the market is about to close before they submit bids. This practice keeps the bids low and rewards those Investigators who are quick rather then promoting the highest value requests.

To overcome this shortcoming, it is possible to use a random closing time, though this approach could adversely affect the outcome if the market closed prematurely. For our experiments, we used the "popcorn" method. In this case, when the market is "popping," bids are coming in and the overall value of the listing is increasing. The market closes when no bids are received over a predetermined period of time.

Another experimental factor that poses a problem is that users do not know a priori how much to bid for a given data take request. Since successful bids cannot be retracted, Investigators have an incentive not to overbid and therefore submit the smallest amount needed to out bid the current request. This situation can produce many small bids and an excessive number of rounds. To overcome this problem, a Vickrey-type auction was used. [4] In a Vickrey Auction, the winning bid "pays" the runner-up price. Thus, if Investigator A submits a bid for 45 points and Investigator B submits one for 60 points, Investigator B "wins" the data take request and is debited 45 points from their account.

Vickrey Auctions provide the incentives for users to be forthright about their bids. If Investigator A tried to underbid by submitting an amount that was lower than what they were willing to spend, Investigator B could submit a bid much higher than Investigator A's and only have to pay the Investigator A's price. Users therefore have the incentive to make bids for the price they are actually willing to pay, which in turn drives the system to a solution faster and reduces the required number of rounds.

Experimental Design

We compared a Serial Draft approach to two market-based approaches (a Simple Market and a Priority Market). A Simple Market is one in which users submit bids with Priority Points. In a Priority Market, users only specify the request's priority. This priority is a measure of the request's importance to the user and has an associated multiplicative factor. The multiplicative factor is applied to the amount of resources requested to determine a bid price for the particular data take request.

Table 3 shows the number of experiments performed for the Serial Draft, Simple Market, and the Priority Market approaches. Results for the

Case	Serial Draft	Simple Market	Priority Market
Simple	Monte Carlo	3	3
Trade-offs	Monte Carlo	3	3
Science Team Simulations			1

Table 3. Number of experimental runs performed for each allocation method and associated case.

Serial Draft approach were obtained by performing Monte Carlo simulations. As the table shows, each approach was done for two types of cases.

The first case was called the "simple" case. In this case, all students had identical preferences (i.e., the same data take requests with the same mission values). This case was designed to study the most heavily conflicted situation where all students desired the same data take time. A second simple case was run where all students had diverse preferences. In this case, all data take requests had different values such that an optimum solution in which each student was able to obtain their high value data takes was possible. This case was done to see if a market-based system could find the optimum solution.

The second case was called the "trade-off" case. This case was designed to see if users would take several lower priority data take requests over their prime request. Trade-off information is nearly impossible to obtain in a serial draft approach. This problem results because trade-off information can only be obtained through detailed questioning of the Science Investigators by the Sequence Integrators. In addition, only the specific questions asked get answered. Sequence Integrators would be hard pressed to ask enough questions to understand the full trade space.

Once trade-offs are made and the market closes, market-based systems move from bidding to the Aftermarket, is a commodities market in which Science Investigators trade any of their resources for any of those owned by another. Aftermarkets are very effective in that both Investigators have to agree to the trade in order for the trade to be completed. This opportunity to barter increases the overall mission value of the timeline.

A final case was performed with Science Investigators to obtain their opinions about a market-based system and its utility to their allocation problem. In essence, the Investigators were used to find out if they could use this approach and whether the approach performed more satisfactorily than one using a serial draft method.

Experimental Results

Figure 1 shows the experimental results when users have identical data take requests. The abscissa axis has the data from the student subjects and a cumulative result. The ordinate axis shows the percent of the mission value obtained using a market-based system as compared to serial draft approach. Thus, a 100% indicates that the same mission value was obtained with a market-based system as compared to the serial draft method. Consequently, a percentage greater than 100% indicates that a market-based system produced a greater mission value for that student.

With identical data take requests, a market-based system was able to produce results comparable to a serial draft method. In addition, for most students, a Priority Market in which they just had to specify a data take priority, did as well or better than when they had to specify a bid price (i.e., a Simple Market). This results from the fact that a Simple Market is less forgiving. Once a bid was accepted, in a Simple Market, it could not be retracted. Thus, a bid with an excessively high number of Priority Points would be accepted and would reduce the student's authority for making subsequent bids.

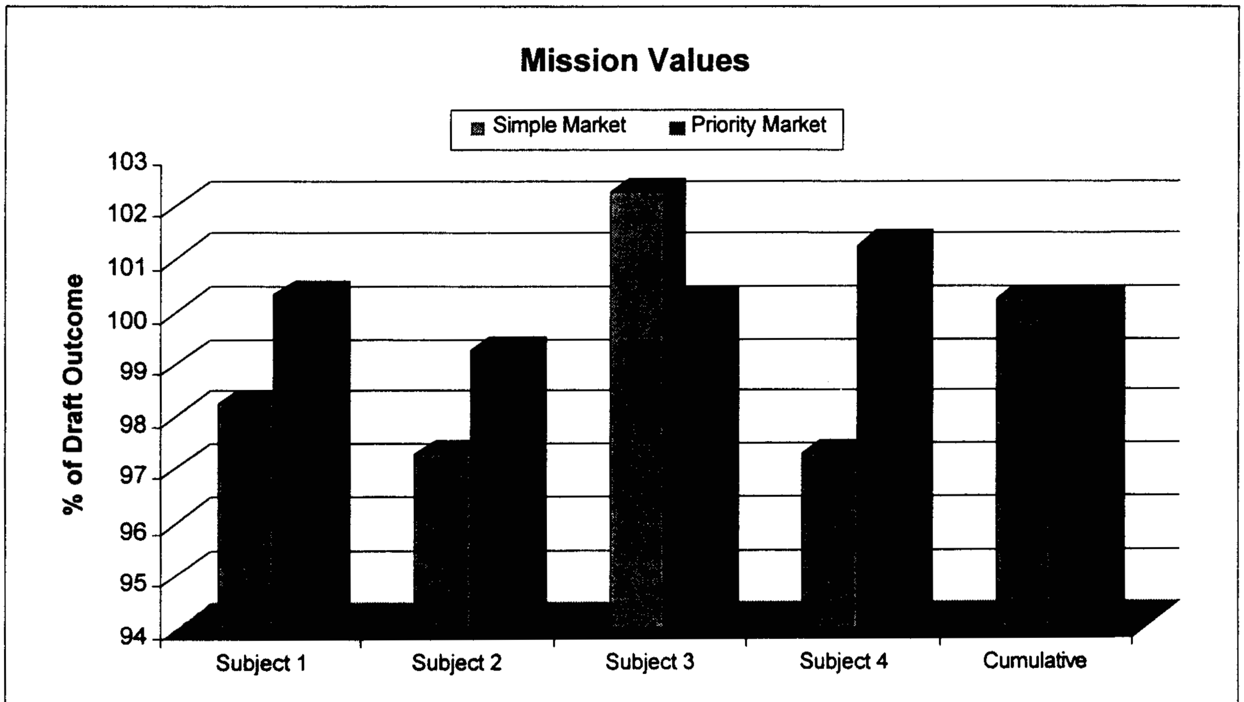

Figure 1. Percent value of market-based approaches compared to a serial draft approach when students have identical data take requests.

In a Priority Market that uses a Vickrey pricing strategy, aggressive bids "paid" the runner-up price. Thus, there was a natural mechanism for preventing excessively priced bids. Only the Priority Points needed to "win" the request were debited from the student's account, thus allowing the individual to use his other remaining points for future bids.

Figure 2 shows the experimental results when users have diverse data take requests. Here again, a market-based system was able to produce results Comparable to a serial draft method. And again, a Priority Market did as well or better than a Simple Market.

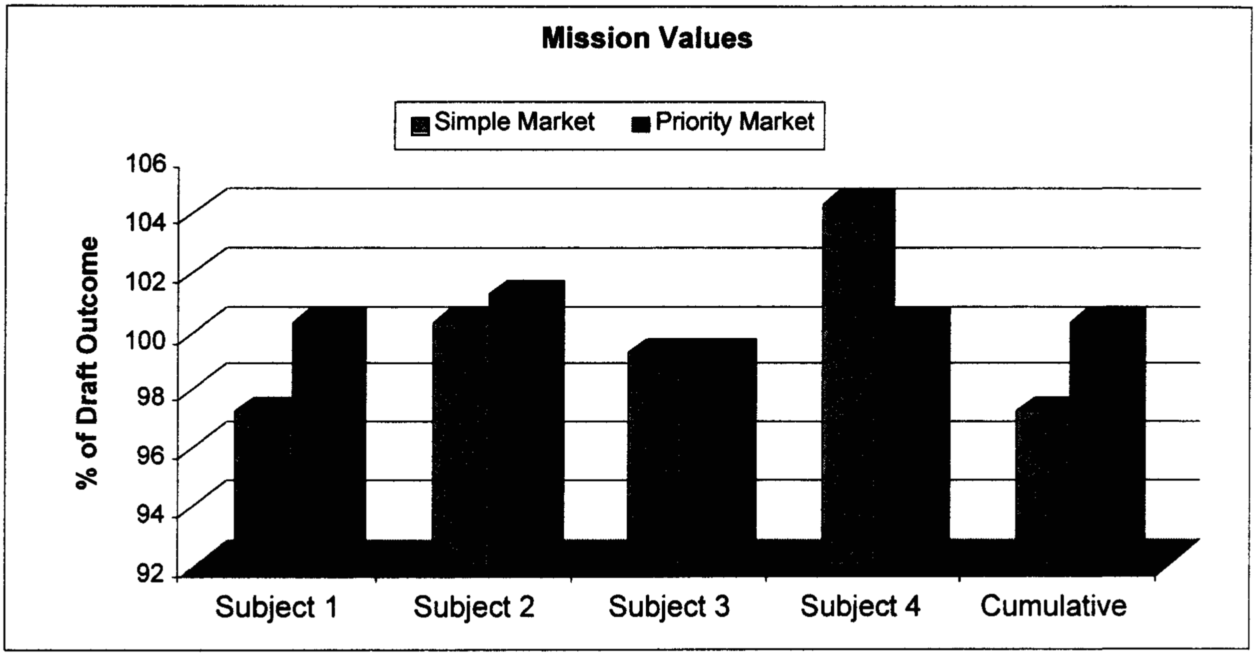

Figure 2. Percent value of market-based approaches compared to a serial draft approach when users have diverse data take requests.

In the Simple Case, both with identical and diverse preferences, results reveal that when few conflicts exist, the market-based approaches (i.e., Simple Market and Priority Market) yield results similar to the serial draft method. That is, market-based approaches were able to find solutions that were comparable to the type of results produced by Sequence Integrators.

Experimental results for the case with "trade-off" information and "trades" are shown in Figure 3. As can be seen in the figure, Investigators had much to gain by making trades. A significant increase in mission value can clearly be realized by selecting a greater number of lower priority data take requests over a few higher priority requests and/or by carrying forward Priority Points. Thus, when there are trade-offs in the number and types of data takes (i.e., when significant conflicts exist) a market-based approach was better than the serial draft approach.

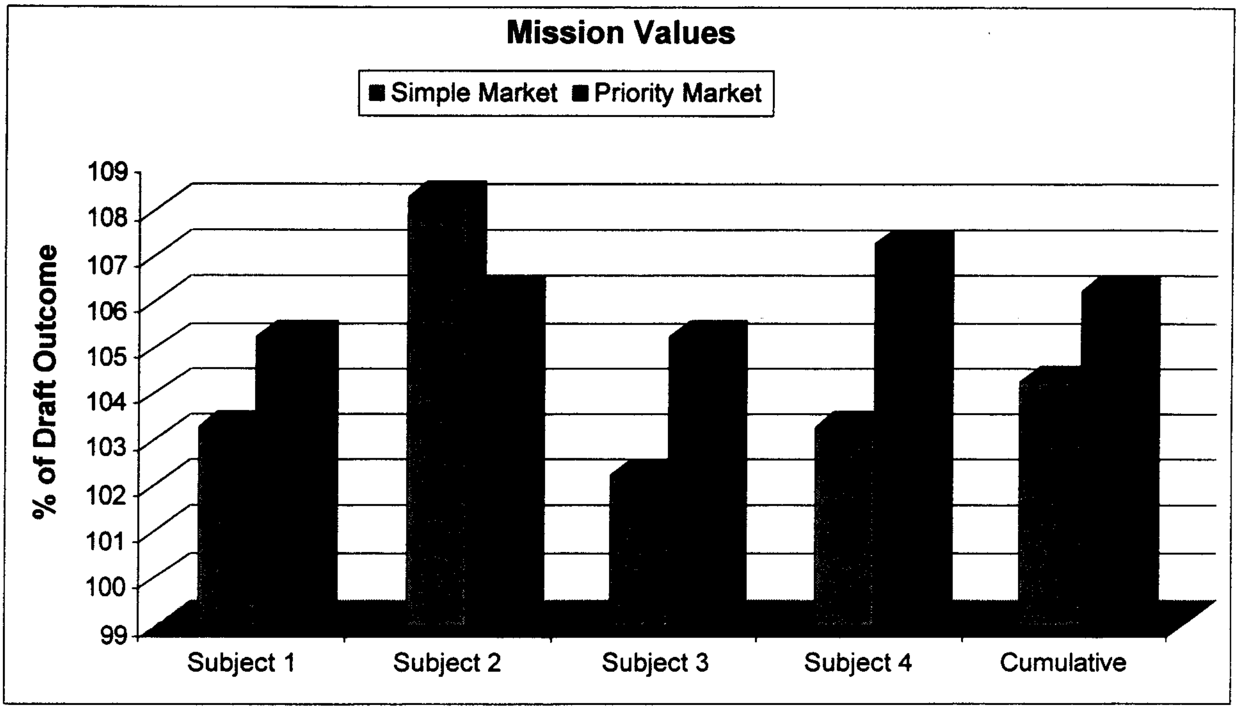

Figure 3. Mission value of market-based approach compared to a serial draft approach when students could take advantage of trade-offs.

The results from this case were superior to a serial draft approach. A Priority Market produced a 2% greater value than a Simple Market method. In addition, a Priority Market converged to a solution in about half the number of rounds needed in the Simple Market approach. This indicates that a market-based system, using a Priority Market, will arrive at a solution faster than a serial draft approach while utilizing fewer individuals to get the same caliber of results.

In addition, a Vickrey Auction Priority Market revealed that specifying a request's priority was more natural to Investigators than specifying a bid "price." The Vickrey pricing strategy made the system more "forgiving" of bids that may have been too high and motivated individuals to submit bids that honestly reflected their true desire for a particular request. Thus, the Priority Market was easier to use, encouraged the generation of accurate bids, and produced the desired conflict-free time-ordered listing in half the time of a Simple Market approach.

Experimental results with Science Investigators from the Jet Propulsion Laboratory revealed that there were few operational problems using a priority market-based approach. There were, however, a number of concerns that were associated more with the experiment than with weaknesses in the market-based system. These concerns included worries over; who determines the initial allocation of points, the questions about whether the experiment was realistic enough (i.e., not enough resources allocated, not enough data takes, etc.), and uncertainties over how long each planning period should be? These issues do not invalidate a market-based system, but accurately reflect the rudimentary capability of the experimental system as compared to one that would actually be used during operations.

With the experiments complete and overall results that indicate a market-based system outperforms a simple ranking approach, the LightSAR Project is moving ahead with the development of a prototype web-based planning tool. This tool will have a realistic Science Investigator interface for submitting data take requests, as well as a market-based solver (with a Vickrey pricing strategy) for developing a conflict-free time-ordered listing that can be directly converted into spacecraft commands for operations.

Prototype Web-Based LightSAR Planning Tool

The following screens are printouts from an electronic prototype mission planning tool being developed for LightSAR. Though the tool is operational, its rudimentary capability reflects the current immaturity of its development, not its full utility to the Project.

Based on experience from past radar missions, the LightSAR Pre-Project recognized the utility of a graphical interface for the input of data take requests. As such, a mercator map projection of the Earth was selected for Investigator input. Figure 4 shows the interface with a low resolution map for testing purposes. Notice that the spacecraft's ground tracks are projected on to the mercator map. Only those ground tracks that occurred during the current planning period were displayed.

Figure 4. Low resolution mercator map of the Earth with spacecraft ground tracks projected on to the map.

To use this tool, a Science Investigator interested in a particular data take request would first define a target region on the Earth. This was accomplished by completing the Data Take Request form (see Figure 5). Notice that a given Investigator had to define an imaging mode, experiment name, look angle, and the latitude and longitude of the target region. Only the experiment name and the default data take priority were optional fields.

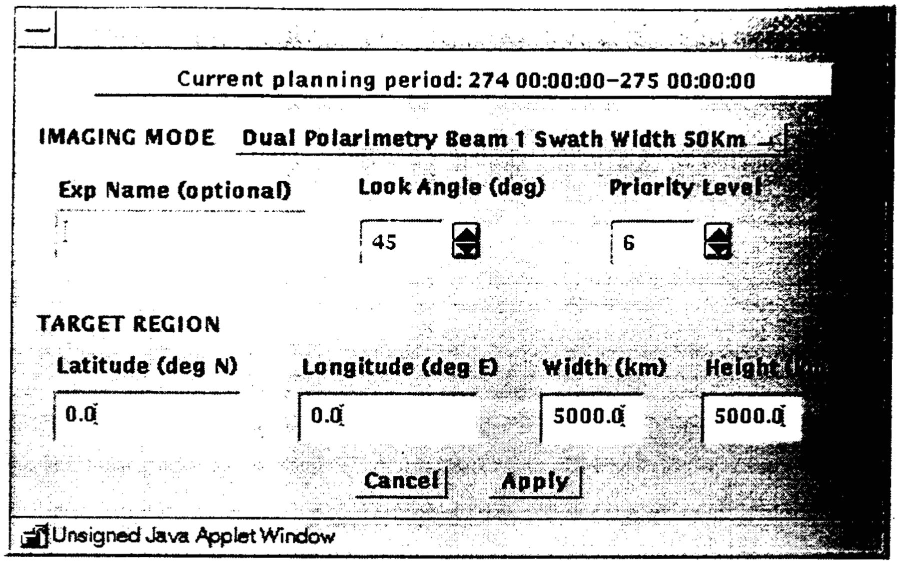

Figure 5. The electronic Data Take Request form.

Once the apply button was selected, the mercator map would show the target region and the associated footprints that crossed the particular region (see Figure 6). These footprints were those associated with the ground tracks displayed on the mercator map.

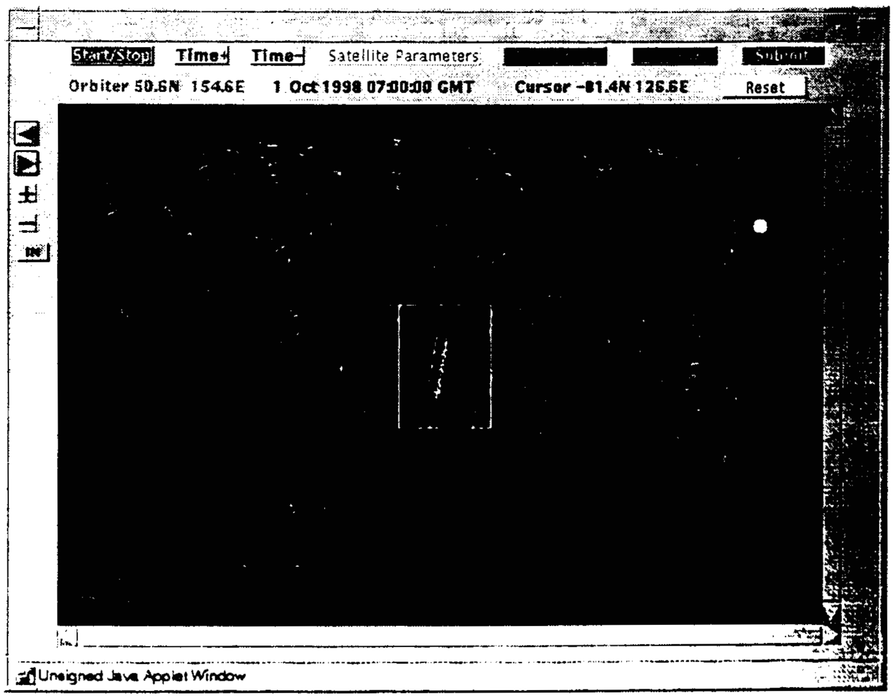

Figure 6. Mercator map of the Earth with the target Region, the associated footprints, and the selected numbered footprints.

Next, the Investigator selected which of the footprints in the target region were of interest. These footprints were selected via mouse entry and numbered according to the time in which they would be collected. Once done, the Investigator selected the submit button.

The interface then returned with the priority selection form. On this form (see Figure 7), the Investigator can see the start and end times for each of his selected footprints and their (defaulted) priorities. The Investigator could than change the priority for any and all of the footprints and use Boolean operators to "and/or" the footprints together. Once the Investigator is satisfied with his request, he selected the "submit" button that transmitted the request to the Solver.

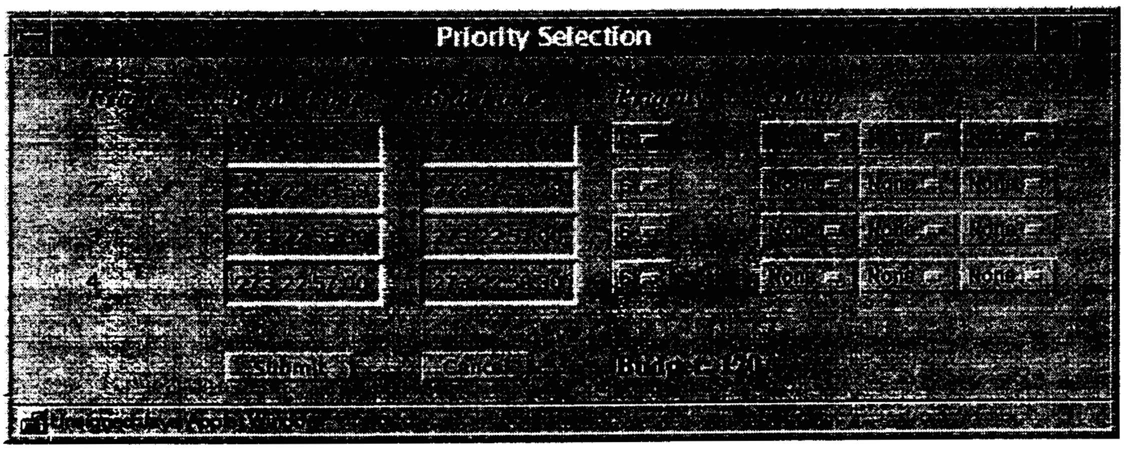

Figure 7. Priority Selection Form with the data take requests.

The Solver returned with a graphical timeline that was conflict-free (see Figure 8). Notice that the timeline is divided into two halves. The top half, in light gray shows the conflict-free timeline. The bottom half, in dark gray, shows those bids that were unsuccessful.

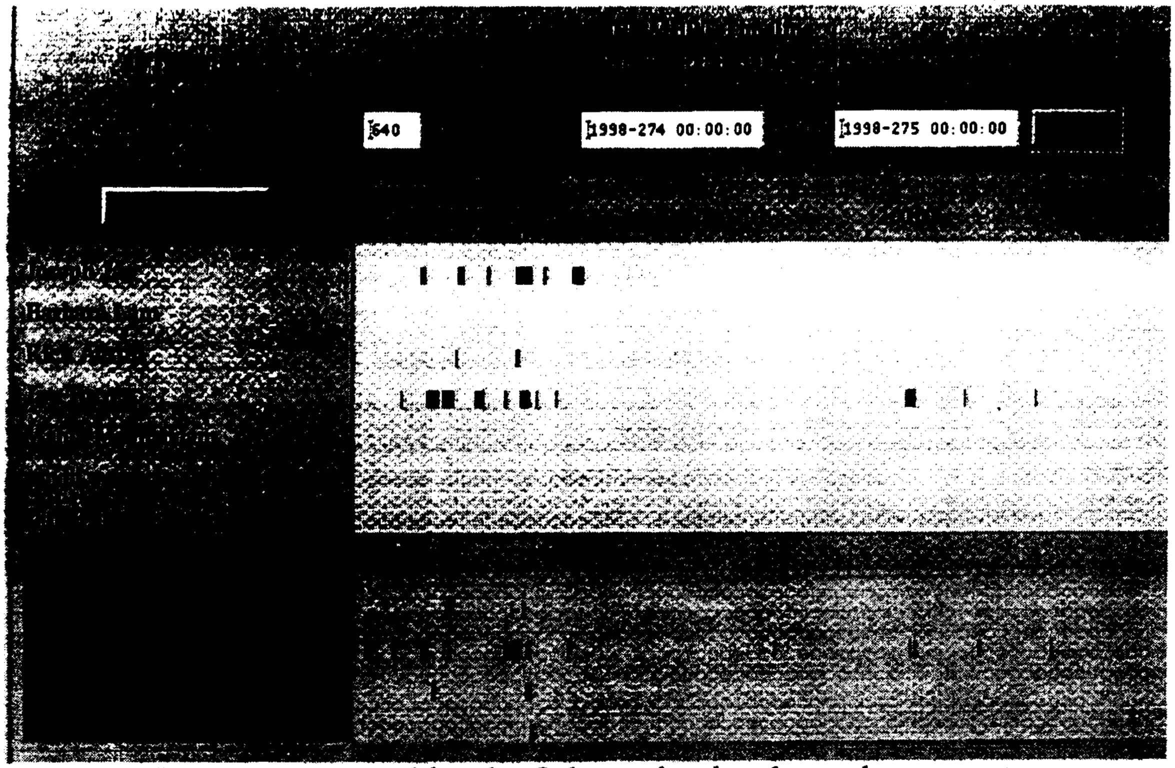

Figure 8. Timeline returned by the Solver of radar data take requests.

From this timeline, Science Investigators could tell if their requests were successful or not. If their bid was unsuccessful, they could either move the time of their request (to avoid the conflict), or increase their bid to try to "win" the time slot.

Once the mission value of the timeline did not increase by 10%, the timeline was "done" and could be ready for the Aftermarket. To do this, the Project would press the "APGEN" button, which would then convert this data to a file that was compatible with the command generation programs used to command the spacecraft. At this point, Science Investigators could trade resources among themselves to increase the value of their data take requests.

Conclusion

Market-based approaches for mission planning outperforms the standard serial draft approach when there are conflicts between numerous data take requests. Furthermore, market-based approaches: 1) provide clearer priority information; 2) remove conflicts, since the timeline is always conflict-free; 3) provide easy access, as the tool is on the Internet; 4) removes the need for Timeline Integration Meetings; and, 5) reduces the number of "appeals" made by Investigators.

Our experiments show that there are no technical issues associated with the operations of a market-based planning tool. Though there are still some management issues to resolve (e.g., who allocates the Priority Points? do Project personnel have the right to veto Investigator trades? how can the graphical interface be made easier to use?), market-based systems have many strengths. They remove the need for Sequence Integrators, can be done faster since all appeals and Integration Meetings are no longer necessary, and can be done remotely from the Investigator's home institution. However, the most important benefit is that that they move the decision-making process back to where the information resides, namely to the Investigators themselves.

Currently, the market-based planning tool is being evaluated further by the LightSAR program. It is envisioned that once the Industry Partner sees the merits of such a system compared to those used on past radar missions, a Web-based, Market-Based system with a user-friendly graphical interface will be the obvious choice when LightSAR launches early in the next millennium.

Key Words

LigntSAR, Mission Planning, market-based systems

Acknowledgment

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

1. TULLOCK, G. (1980). "Efficient Rent Seeking", *Toward a Theory of the Rent-Seeking Society,* edited by James Buchanan, Robert Tollison and Gordon Tullock. College Station: Texas A&M Univ. Press.

2. WESSEN, R. R. & D. PORTER (1998). "Market-Based Approaches for Controlling Space Mission Costs: The Cassini Resource Exchange", *Journal of Reduced Mission Operations Costs,* Vol. 1, No. 1, pgs 9 - 25.

3. WESSEN, R. R. & D. PORTER (1999). "A Market-Based Approach for Manifesting Shuttle Secondary Payloads", *Journal of Spacecraft and Rockets,* Vol. 36, No. 1, pgs 142 - 147.

4. VICKREY, W. (1961). "Competitive Sealed Bid Tenders," *Journal of Finance,* vol. 16, pp. 8 - 37.